Towards Predictable Low Latency DOCSIS Services

A Technical Paper prepared for SCTE by

Dan Rice
VP Access Architecture and Technology, CONNECT
Comcast
4100 E Dry Creek Road, Centennial CO
daniel_rice4@comcast.com

Sebnem Ozer, Ph.D.
Senior Principal Architect, CONNECT
Comcast
sebnem_ozer@comcast.com

James Martin, Ph.D.
Professor Emeritus
Clemson University
Jmarty@clemson.edu

Table of Contents

List of Figures

1. Introduction

The Internet is based on fundamental design goals that emphasize simplicity and reliability. Early on, the idea was that instead of a centrally controlled network that was a gatekeeper on new features (such as telephony 'star' codes) most of the intelligence was decentralized into the edge systems to help enable a tremendous wave of permissionless innovation at the application layer. Any complex network-oriented processing was to be performed by the host devices. A fundamental difference of today's Internet is 'predictable service qualities', which can be achieved only by 'smart' components in an end-to-end path. Over the years, applications have adapted as well. A classic success story is Internet streaming. Through maintaining an appropriate size playback buffer and aided with adaptive bitrate control, HTTP-based adaptive streaming (and UDP equivalents) is the dominant application.

As society moves from the information age to the age of the M2M communications, we anticipate further significant changes to the Internet as well as significantly different applications. Machines are making decisions based on data, many times in real-time. There are wired and wireless scenarios to consider. Many of these applications have yet to be invented, but based on early examples such as cloud gaming, enhanced videoconferencing, coordinated autonomous vehicles, drone swarms working in a coordinated manner to carry out missions, multiuser virtual-reality gaming, the Internet needs to be re-examined to determine what must be addressed so that technological breakthroughs in devices and systems are not held-back by Internet performance limitations. These concerns have motivated the key standards bodies to pursue broad initiatives that will enhance their respective technology's ability to support emerging applications systems that require predictable service qualities.

Cable operators are evaluating Low Latency DOCSIS (LLD) to determine how the new technology might improve subscriber's perceived quality and how an operator might leverage the technology for new services [1, 2, 3].

In this paper, we first present LLD lab analysis that shows promising results. We then discuss end-to-end factors that are critical in providing predictable end-to-end LL services. We conclude our paper with an architecture including network and service components to deploy an effective LL system.

2. Low Latency DOCSIS: Promising Results

The term Low Latency DOCSIS (LLD) reflects the general Internet and cable community's efforts to offer an architecture that satisfies the requirements for both non-queue building (NQB) traffic and queue-building (QB) traffic over the same physical network [4, 5, 6]. The core ideas behind LLD are novel and show promising results based on preliminary studies and trials. The goal is to enhance the end-user experience and to allow operators or over-the-top service providers to develop new services.

LLD is compatible and aligned with a new generation of Internet flow control, with Low queuing Latency, Low Loss, and Scalable throughput (L4S) technology [4] which supports end-to-end latency that manages queueing delays with congestion notifications that will result in flow rate adjustments, thus avoiding the packet drops and retransmission as much as possible for the NQB flows. The goal is to transition to substantially lower queuing delays across the network while coexisting with 'classic' congestion controls used widely in the Internet today. This will become important not only to the development of new applications as described here, but also to improve customer experience as sufficient network speeds are now becoming widely available to all Internet users and higher speeds may no longer improve the Internet experience, especially for residential consumers.

LLD aims to support non-queue-building applications that send data quickly and don't cause latency in the presence of queue-building applications from and to the same subscriber's modem. The main components are coupled dual-queue with weighted scheduler, proactive grant scheduling and optimized MAP timing and channel settings [6]. The aim is to support ~10ms DOCSIS round-trip-time (RTT) latency for the 99th percentile of LL flows, a ~10X improvement to today's DOCSIS deployments.

The initial tests show promising results. The following upstream examples are discussed to describe the benefits of LLD in the presence of high home network utilization. The lab tests are carried over an iCMTS with deployment configuration and D3.1 CMs supporting LLD features. A simplified version of the setup is shown in Figure 1.

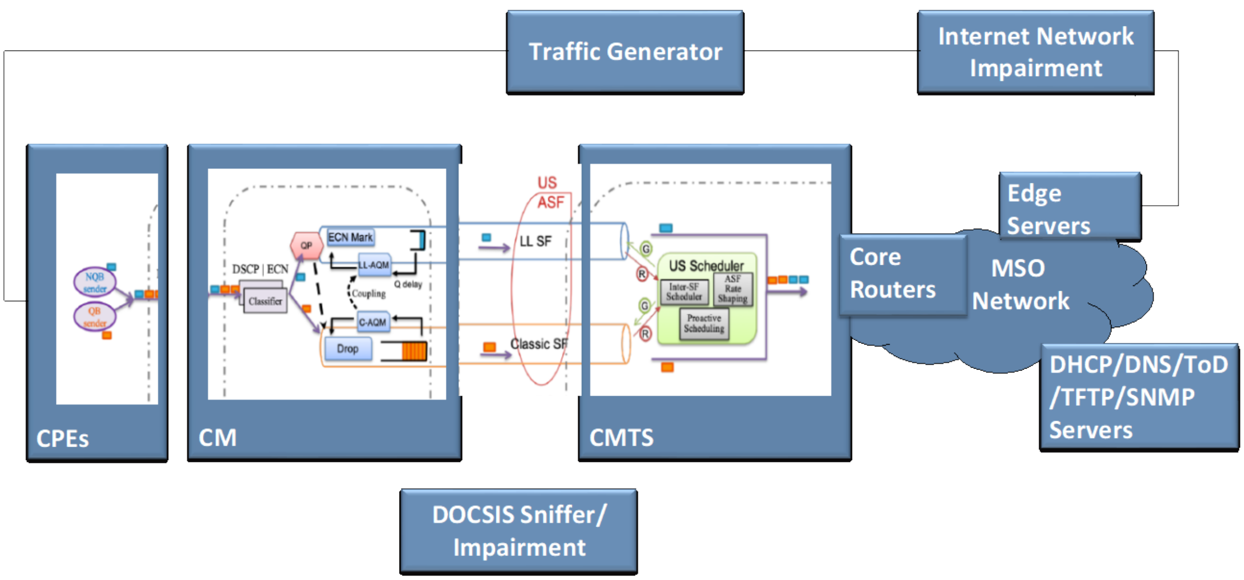

Figure 1 – LLD Lab Setup

The queuing and media access are the major DOCSIS latency sources while propagation, serialization, encoding, and switching also contribute to the latency. Figure 2 displays example idle latency values for a VoIP type traffic in the upstream direction. Idle latency refers to latency when there are no other simultaneous home traffic flows competing with the test traffic for the shared resources. Delays due to propagation, FEC, ATDMA Interleaving/OFDMA Framing, CM/CMTS MAP processing, queuing (~MAP timing/2) and other processing are computed for a total latency estimation as shown in blue line in Figure 2. A traffic generator is used to create VoIP type traffic and measure latency as shown with circles in the same figure. As expected, propagation latency, described by the data point annotation, is the major factor depending on the distances between the CM and the US scheduler as a function of CMTS. Today, most headend-to-CM distances are smaller than 160 km, deployed with typical DOCSIS timings, which provides good latency performance. If better latency bounds are targeted, US scheduler can be deployed closer to the subscriber as a microservice or proactive grant scheduling techniques can be integrated.

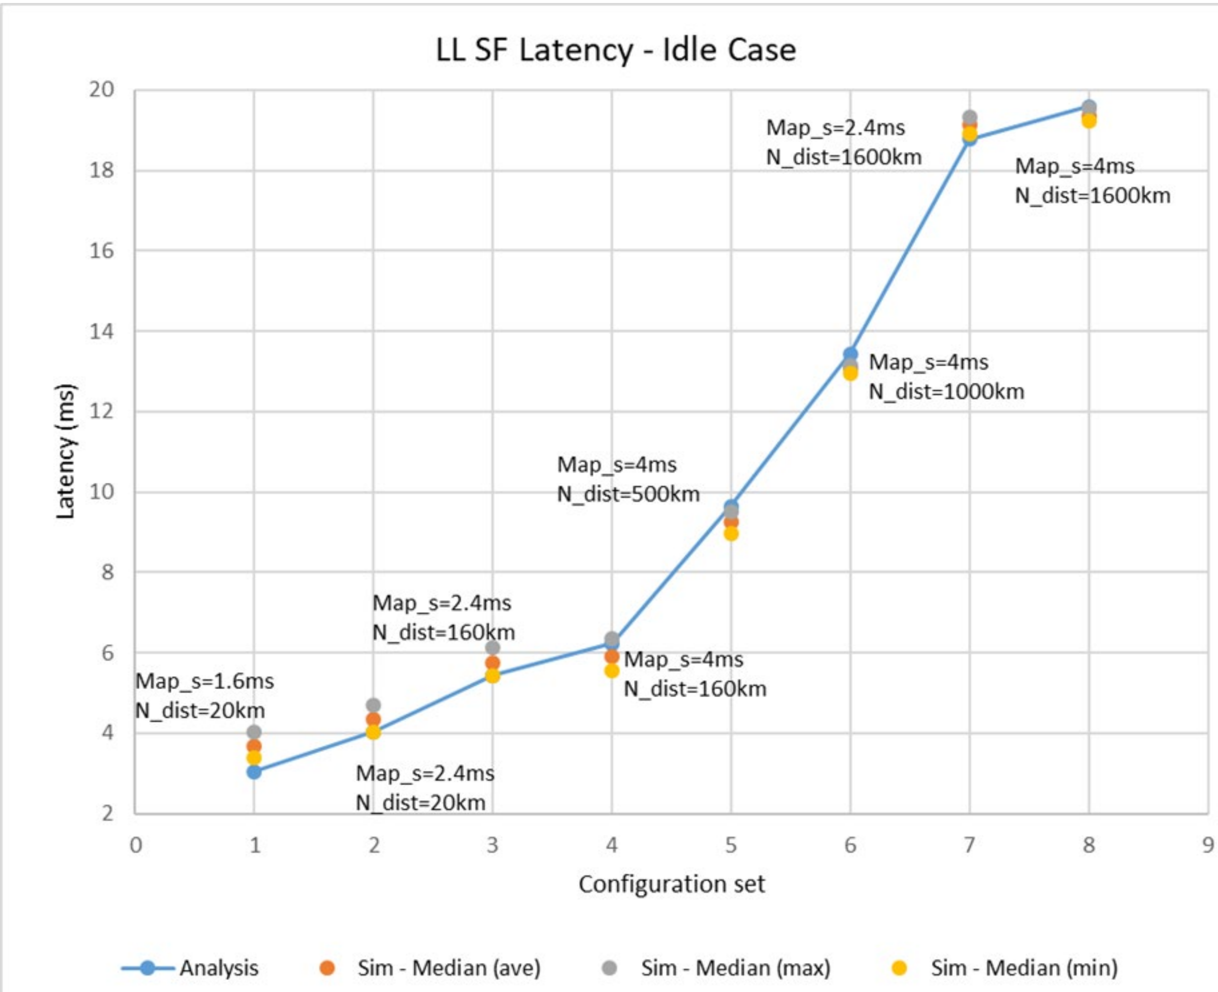

Figure 2 – Idle Latency Analysis

The idle latency provides important information on the achievable minimum latency and typical values when LL traffic is not affected by other traffic in the CM. Today in many homes there are many users or sources of traffic that use the network simultaneously with work and school from home that increased dramatically in recent years. As a result, while setting a lower bound on latency performance, is a less effective metric for user experience. However, LL services require consistent performance values such as jitter.

Latency under load or working latency computed during high home network congestion levels provide upper bounds for latency and jitter a LL service may be subject to, depending on the speed tier rate and conditions of the shared resources. Figure 3 displays an example use case where two audio flows are generated for the same CM, one marked with LL DSCP value and the other one marked with DSCP 0x00. The dual queue LLD features are tested. The top graph displays the latency over time (green line) for unmarked audio flow while the middle graph displays the latency over time for an audio flow marked per LLD specifications. The concurrent TCP flows ‘throughput, representing other traffic flows in the home, is displayed in the bottom graph. As the TCP flows fill up the speed tier bandwidth, both audio flows’ latency values are affected proportionally to TCP flows’ throughput fluctuations as seen in the figure. However, the unmarked audio flow’s latency increases proportional to the TCP traffic RTT as it shares the same queuing and scheduling weight while the marked audio flow’s latency is bounded with a ~6X improvement. Not only is the absolute delay much lower, but the variation in delay is within a much smaller window which is a critical key improvement for most applications. Note that the MAP time and other physical layer channel settings are not yet optimized for these results set, and will lower latency even further.

Figure 3 – US LLD example showing the impact of QB traffic on NQB traffic latency

In Figure 4, two marked NQB traffic flows share the same LL queue while QB TCP traffic and unmarked NQB traffic is transmitted through the classic queue for the dual queue approach. In [5], suggested unresponsive NQB traffic load is 1 Mbps or a couple of packets per RTT. The marked NQB flows share the resources fairly with bounded latency results (top two graphs), while unmarked NQB traffic (third graph) has inconsistent latency fluctuating depending on the TCP traffic (bottom graph). The latency histograms and Complement Cumulative Distribution Function (CCDF) are shown in Figure 5, with top two graphs for marked NQB flows, third graph for unmarked NQB flow and last graph for CCDF of all three NQB flows. In this example, there were a few packets with latency close to 19 ms and 99%ile of packets have latency smaller than 18 ms. Further optimization of LLD feature parameters and DOCSIS settings can reduce this latency range further. For the same settings, the latency improvement is close to 5X for 99%ile range. The flows in this simulation are longer than those of Figure 3 where TCP flows are ramping up for most of the simulation time. The latency variation for the unmarked NQB traffic can still be observed in Figure 4 during the steady-state TCP conditions as well.

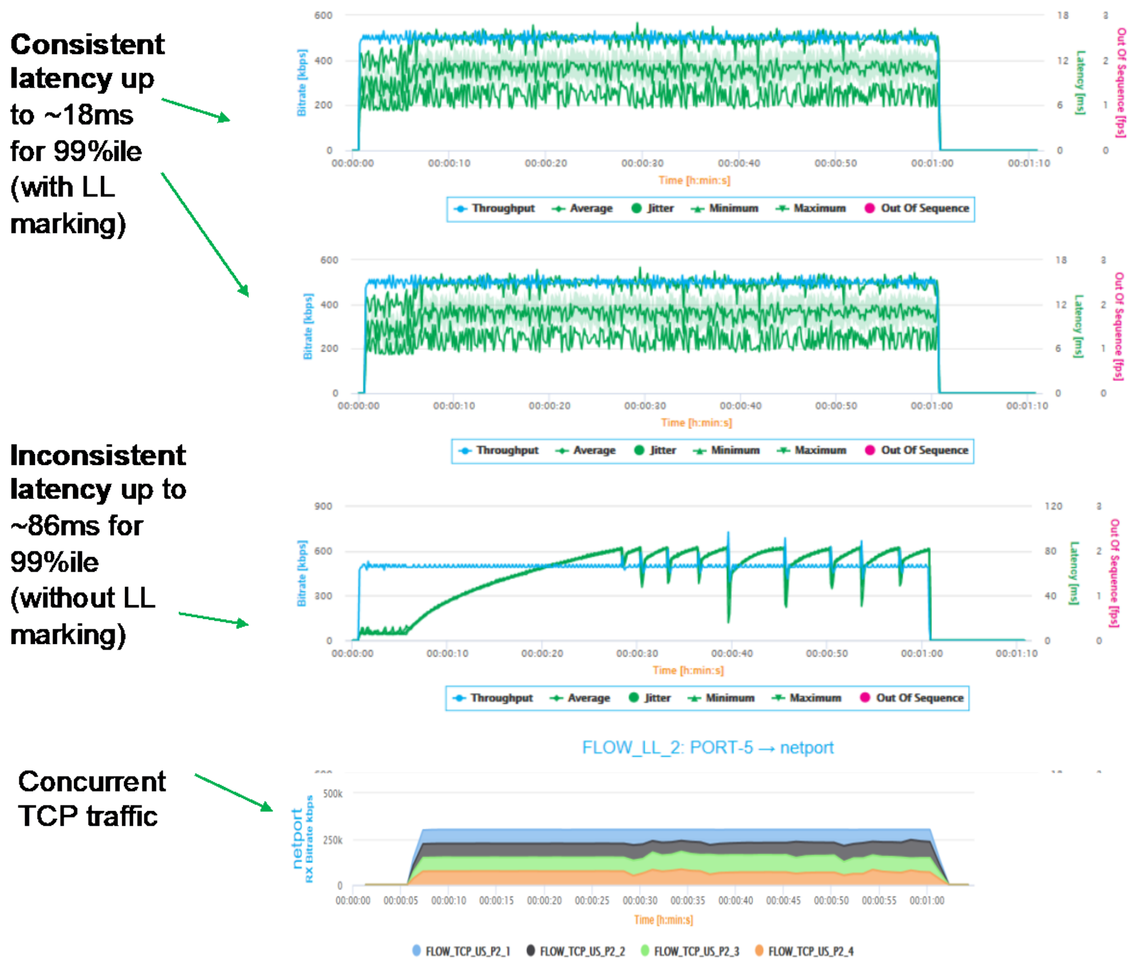

Figure 4 – US LLD example with two marked NQB traffic flows of the same rates in the CM

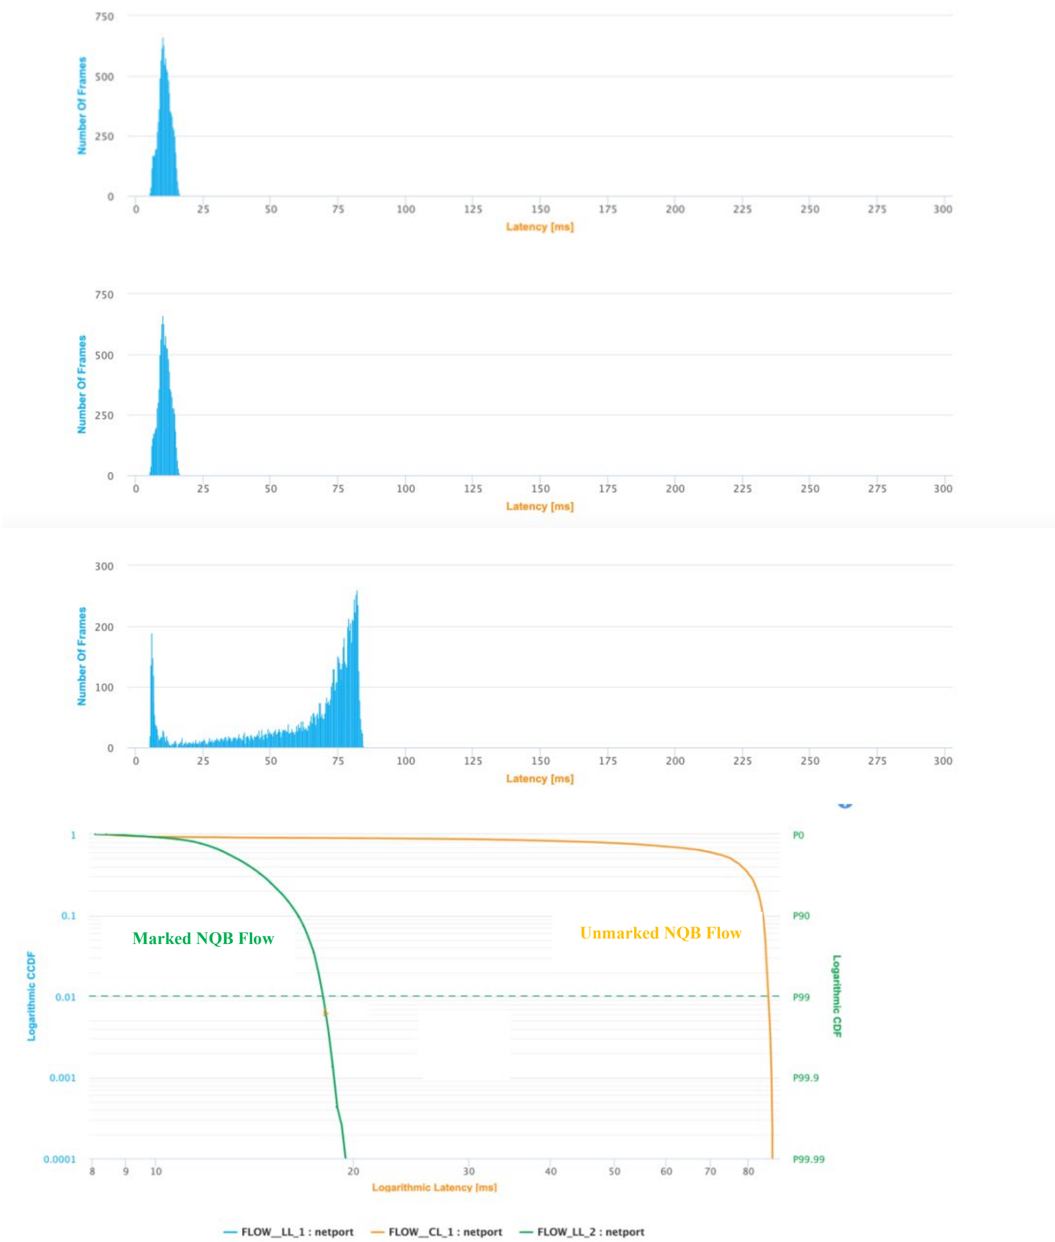

Figure 5 – Latency statististics for US LLD example with two marked NQB traffic flows of the same rates in the CM

A similar use case with two marked NQB traffic flows of different rates are displayed in Figure 6 and Figure 7 to show that the flows are served fairly in terms of latency, jitter and throughput. While improving the 99th percentile latency by ~6x.

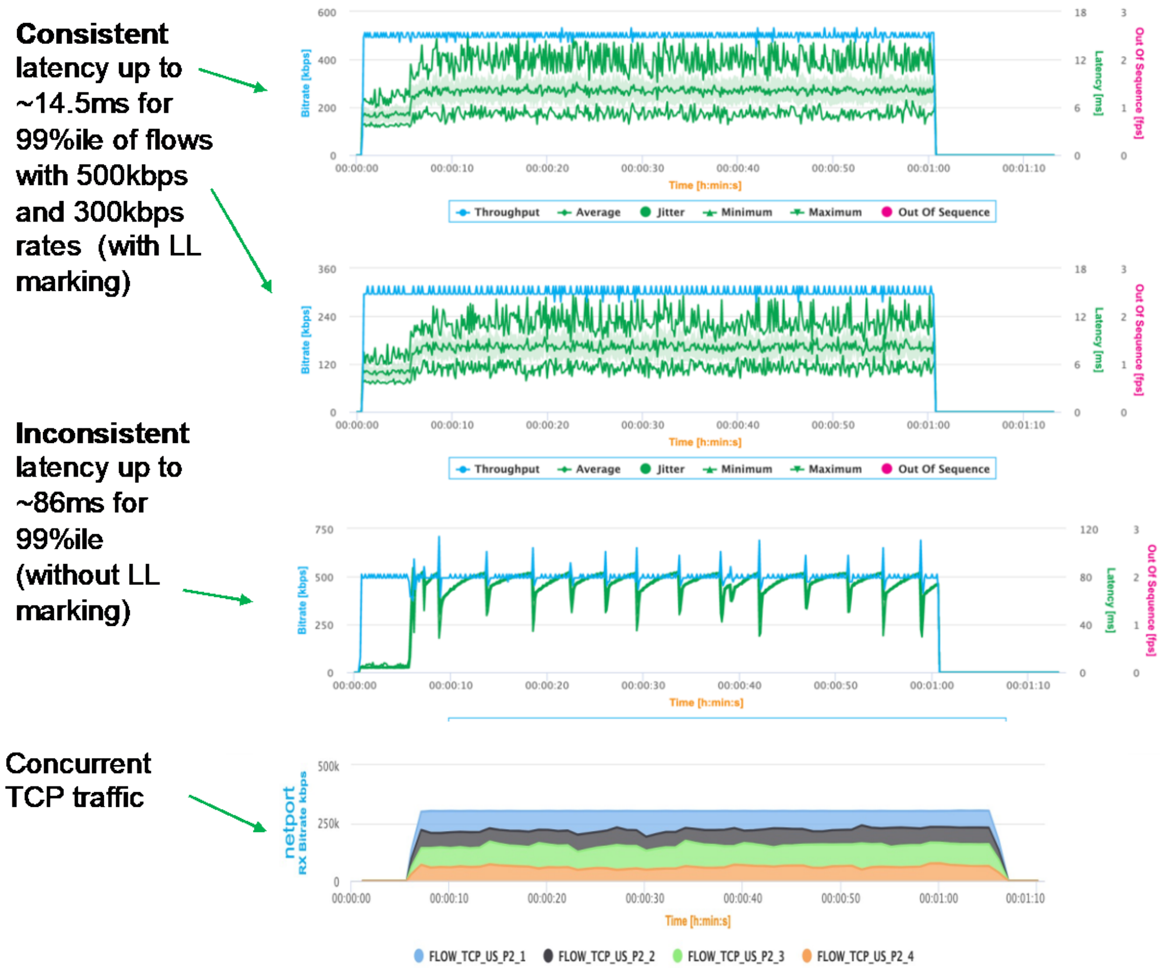

Figure 6 – US LLD example with two marked NQB traffic flows of different rates in the CM

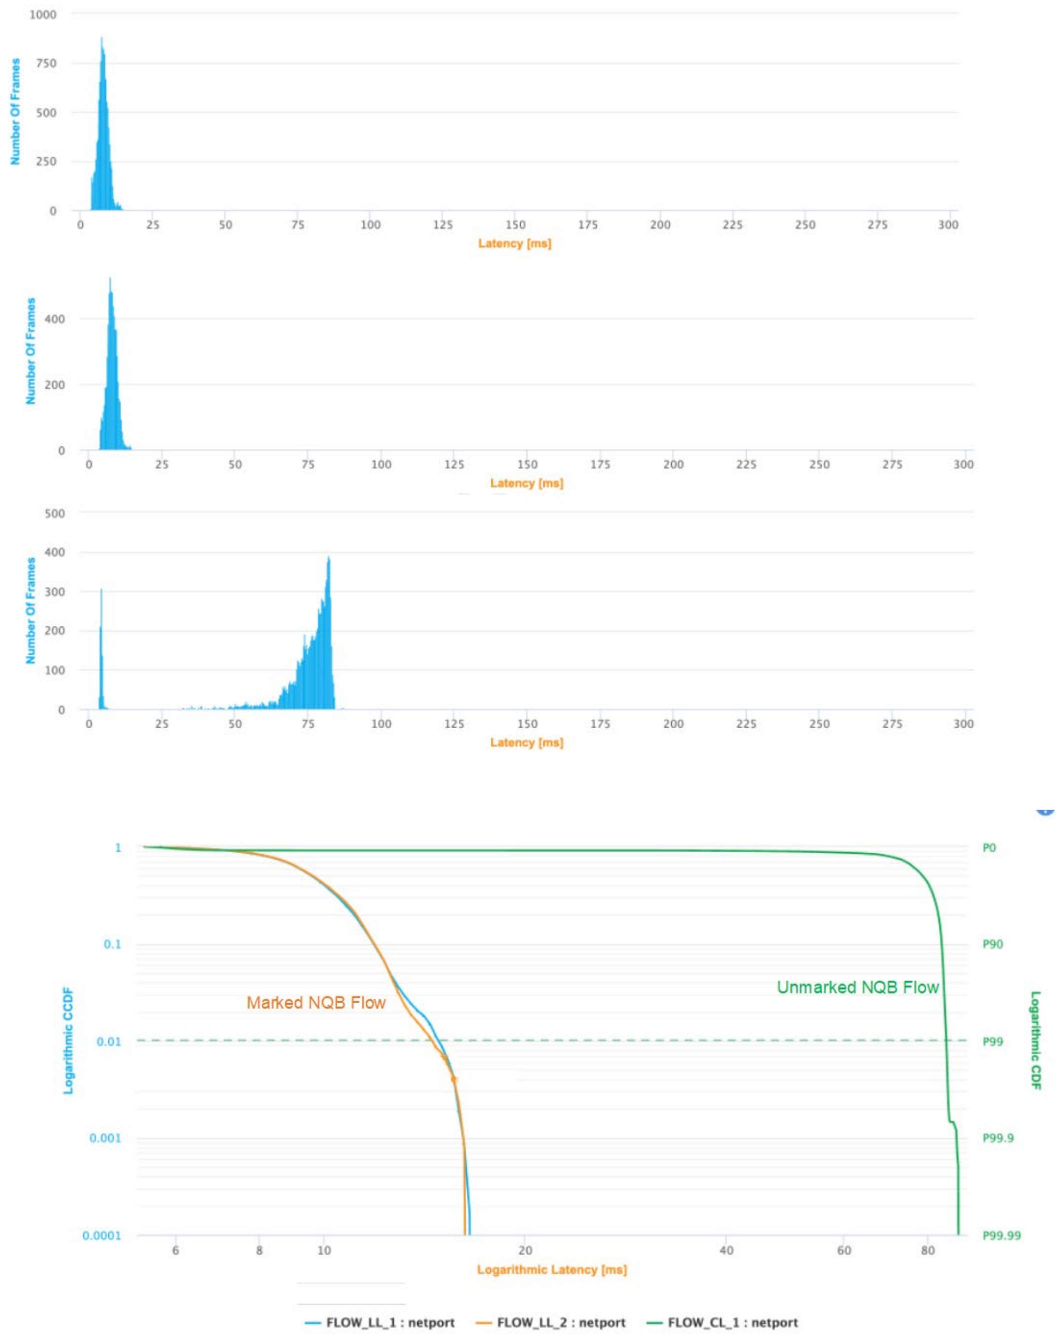

Figure 7 – Latency statististics for US LLD example with two marked NQB traffic flows of different rates in the CM

For applications requiring tighter latency bounds or less variations due to use cases such as serving group congestions, Proactive Grant Scheduling as defined in [6] may be preferred. Although PGS eliminates the latency due to request-grant cycle in the upstream direction, it may create network inefficiency if the grants timing and NQB traffic characteristics are not matched well. In this case, unused grants would decrease network efficiency while less grant allocations might increase the latency. However, as shown in Figure 8, the 99th percentile latency and jitter are further improved by and incremental ~65% with this technique increasing the overall improvement to > 13X. Cable operators may apply PGS for certain services or extend the implementation with traffic activity detection.

Figure 8 – Latency statististics for US LLD example with PGS

3. Predictable End-To-End Latency

As discussed in the previous section, LLD improves the latency and jitter for the NQB traffic. Since NQB traffic doesn't share the same queue as QB traffic, packet loss due to aggressive QB traffic is avoided. While DOCSIS network segment can benefit from LL features, the subscriber's QoE depends on the end-to-end latency. In this section, different factors that cable operators must consider for a predictable latency are discussed. Although latency is mentioned as the main metric, in fact, jitter, packet loss and throughput are analyzed concurrently. Reliability, security and service availability are other metrics that operators consider for a complete service assurance.

3.1. E2E Marking

LLD and LL implementations in other network segments depend on marking unresponsive NQB and L4S traffic with LL DSCP and ECT(1) [4,5]. In traditional systems so far, operators have bleached or remarked marked packets entering their network. Wi-Fi gateways and routers in the upstream direction and interconnect routers and CMTSs in the downstream direction bleach or remark these packets.

Operators must assure their network components do not bleach or remark (except CE marking as explained in [4] and DSCP conversions as explained in [5]) LL packets in their network. Queue protection schemes [4,6] and new security features may be applied to detect NQB traffic violation. Negotiations with peering partners and edge servers may be used to avoid the issue with network components in the segments outside of operator's control.

Although, techniques such as mirroring the marking in the downstream per the received marked packet in the upstream may be used, symmetry is not always guaranteed. Application detection should be used only if subscriber has full control and/or a paid service is deployed.

3.2. End-to-End Measurements

Cable operators have been deploying latency, jitter, packet loss and throughput/speed measurements in their network for idle, LUL/working and real customer traffic [1]. Techniques that can measure latency in the access network as well as latency outside of operator's network has been also deployed [1]. These platforms must be upgraded to measure these metrics per LL marked and classic packets and flows.

These measurement techniques must be analyzed with other monitoring features such as utilization levels. device conditions, routing options, channel conditions and many more resource conditions that can affect e2e latency. As latency targets get lower and/or speed tier rates increase, traditional monitoring windows may not be adequate for this analysis. For example, 5 min averages for channel utilization have been widely used in cable operators' networks.

However, bursts (even microbursts) at smaller windows may have a big impact not only on flow control for very high speed tiers, but also on LL services. Traditional monitoring settings may not catch the correlations and causations. An example is shown in Figure 9 for a serving group with highly variable utilization bursts. 5 min values correspond to 98% of utilization measured within 15 sec and reported every 5 min. 15 sec values correspond to 98% of utilization measured and reported every 15 sec. It can be seen that 15 sec values report high burst rates while 5 min values are smoother. The same behavior may be observed for home network utilization and other channel and network conditions.

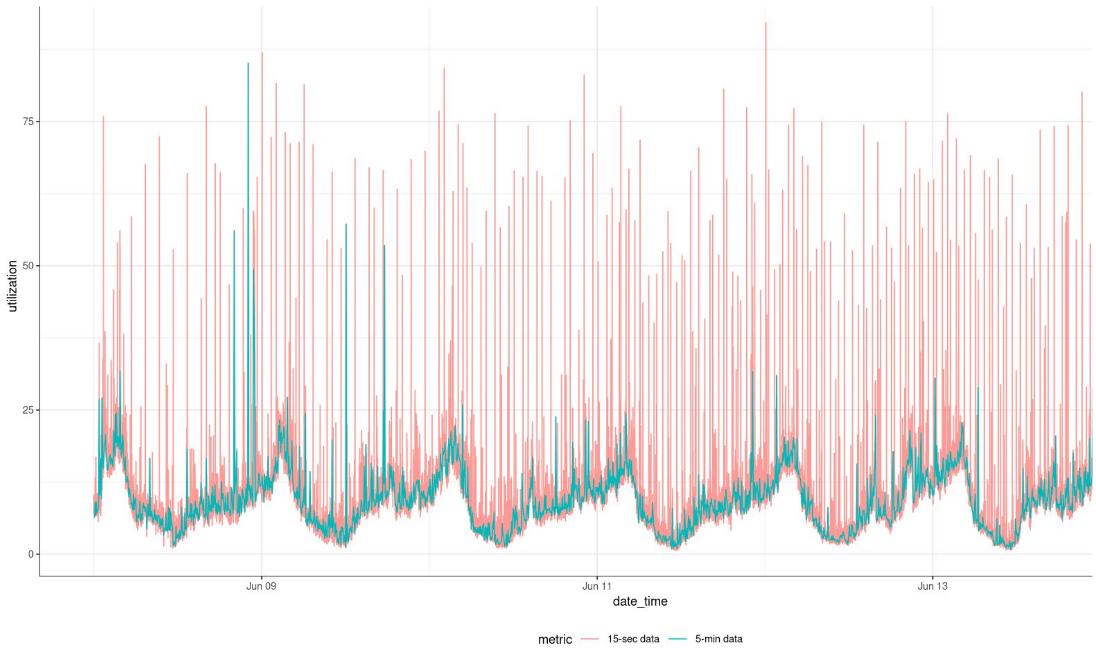

Figure 9 – Utilization statistics reported with 15 sec and 5 min windows

LL services QoE may correspond to different QoS levels. Although it is not a straightforward process, operators developing platforms to map QoS levels to LL service QoE will be able to program their networks in an optimized way.

3.3. End-to-End Optimizations

A cable operator's network may be approximated as a tandem queuing and servicing. As shown in the literature, in tandem queuing, the arrival rate and distribution in a network segment depends on the previous network's arrival and service processes. Therefore, analyzing the network segments with interdependency can help the operator to estimate the end-to-end latency and find ways to optimize the network features in an end-to-end way.

Mathematical models for complex tandem queuing system with general arrival and service processes may not exist. Operators may use measurement and machine learning techniques to detect the dependencies of network segments.

Edge servers may be deployed for cases where dependency outside of the operator's network creates an obstacle for certain LL services, depending on the cost and latency analysis.

Another factor to consider is how marking is passed based on the next segment's LL support. For example [5] proposes remarking LL DSCP to a lower DSCP value if the next network segment is not ready to support LL services and a higher DSCP may create unfairness or starvation for other services.

3.4. Network Segments with Classic ECN

Although a rare use case, network bottlenecks employing a shared queue that implements an AQM algorithm that provides Explicit Congestion Notification signaling according to RFC3168 may create issues for the L4S deployment [7]. ECT(1) usage for L4S traffic classification are defined in RFC8311 to update RFC3168. If end-to-end path includes network segments with RFC3168 AQMs, having L4S compliant congestion controlled flows instead of all classic congestion controlled flows may cause a lower throughput of classic congestion controlled flows in the same queue. As explained in [7], this problem is rare and significant only if RTTs are long and rates are high for long running flows. [7] proposes preferred and other options, including using edge servers, upgrading AQM implementations, fairness improvement techniques etc. Operators may use end-to-end monitoring and QoE assessment techniques to discover such cases and deploy options defined in [7]. IETF specifications have recommendations for other cases that may be incompatible for L4S deployments during transitioning to a larger ecosystem support.

3.5. Standards Compliance and Larger Ecosystem Support

Although IETF and CableLabs specifications are the main standards cable operators rely on for LL feature implementations, most cable operators support other technologies in their networks, such as PON, 4G/5G, Wi-Fi and wireless IOT technologies. For a predictable E2E latency values, the standards' LL approaches such as marking and congestion notification techniques must converge. For example, IEEE 802.11be (Wi-Fi 7) specifications introduces new Access Categories and mapping for LL services, that comply with LLD features.

In addition, applications (e.g. online and cloud gaming, videoconferencing, interactive real-time streaming, VR/AR), OS (e.g. Windows, MAC, Linux) and API protocols (e.g. Webrtc, Element) providers must also support LL marking, accurate congestion notification and scalable congestion control techniques. LL services may be supported seamlessly and in a predictable way with larger ecosystem support as every party will win with a common goal.

At the IETF 114, sponsored by Comcast, and held in Philadelphia the week beginning July 23, a Hackathon to test interoperability of L4S, ECN and LLD took place.[9]. There were over 30 participants including cable operators and large Internet application and technology providers including Google, Apple, Meta, Netflix and NVIDIA among others. During this hackathon network infrastructure for DOCSIS, Wi-Fi and 5G were all tested. Results for LLD traffic with an AppleQu8ic implementation are summarized in Table 1.

Table 1: IETF 114 Hackathon results with very compelling improvements in Latency and Jitter.

Upstream	Classic AQM Network	L4S + LLD Network
Upstream Flows	P99 – 30 msec P99.9 - 125 msec	P99 – 9 msec P99.9 - 10 msec
Downstream Flows	P99 – 56 msec P99.9 - 96 msec	P99 – 1.1 msec P99.9 – 7.8 msec

4. Conclusion

LLD promises predictable low latency service support if end-to-end architecture and deployment strategies are well defined. It requires many service and network components to interact in a harmonious way as shown in Figure 10. To support such an architecture requires disruptive changes in traditional cable operators' networks. Operators should support Software Defined Networking (SDN), Network Functionality Virtualization (NFV), Data-driven automation and service agility for the best LL service support.

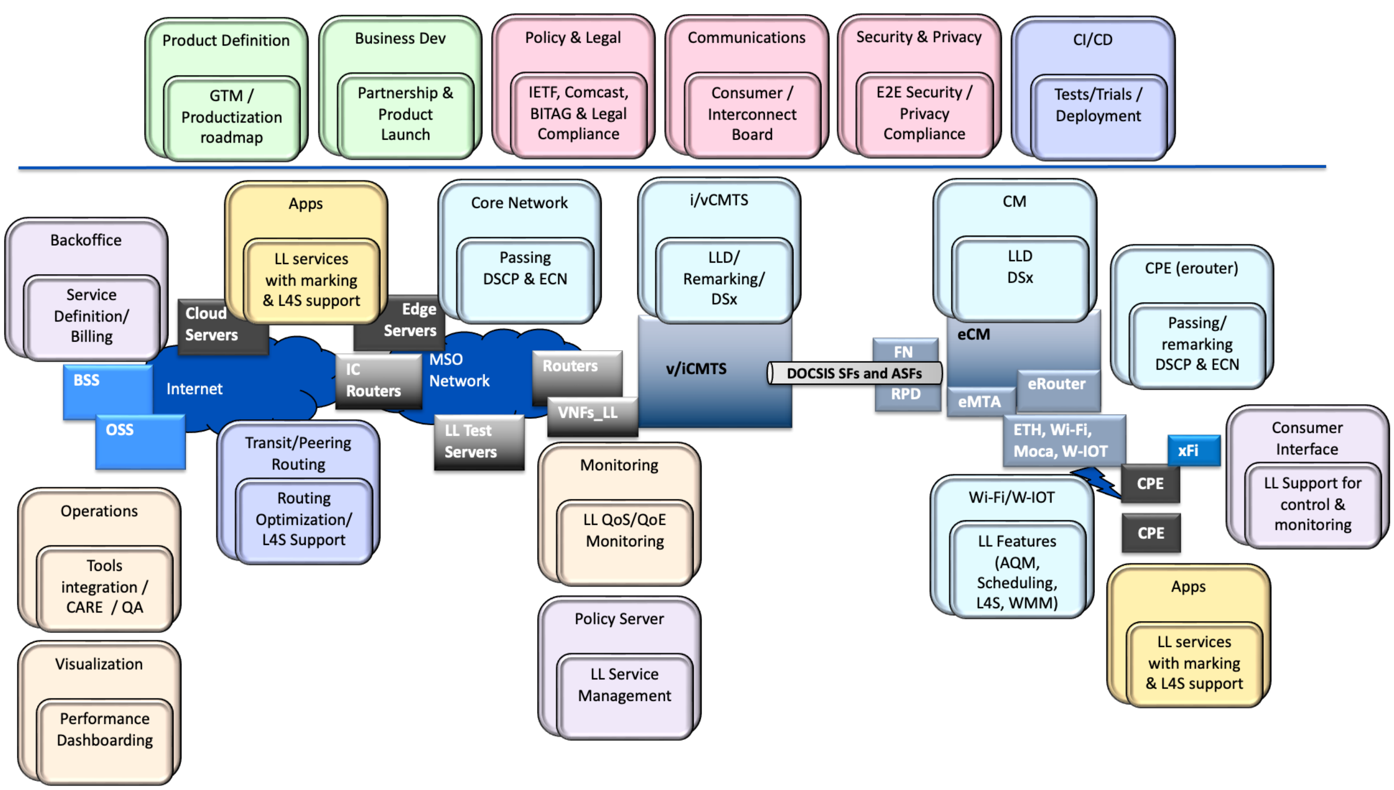

Figure 10 – LL Service and Network Architecture

Acknowledgments

The authors would like to thank Joe McHale, Sumi Chandrashekar, Garrett Miller, Chad Humble for their LLD lab support and Maher Harb for his utilization statistics analysis.

Abbreviations

API	Application Programming Interface
BSS	Business Support System
CM	Cable Modem
CMTS	Cable Modem Termination System
DHCP	Dynamic Host Configuration Protocol
DNS	Domain Name System
eMTA	Embedded Media Terminal Adaptor
HFC	Hybrid Fiber Coaxial
HSD	High Speed Data
MAC	Medium Access Control
MSO	Multiple System Operators
OSS	Operations Support System
PON	Passive Optical Network
QoS	Quality of Service
TFTP	Trivial File Transfer Protocol
ToD	Time of Day

Bibliography & References

[1] *Fastest Path to Low Latency Services: How Can Cable Operators Deliver Consistent Latency by Following an Efficient and Future-Proof Design Path?,* by S. Ozer, A. Tunstall, C. Klatsky, D. Rice, J. Livingood, J. Chrostowski, J. Raezer, J. Gerson, M. Dolley, P. Sarathy, S. Subbaraj, S. Cho, T. Graffa SCTE Expo 2021

[2] *Configuring and Deploying Low Latency DOCSIS Networks*, by Greg White & Karthik Sundaresan, SCTE Expo 2021

[3] *Low Latency Docsis: Concepts and Experiments*, by Tushar Mathur, Ram Ranganathan, Greg Gohman, Bob Zhang, SCTE Expo 2020

[4] https://datatracker.ietf.org/doc/draft-ietf-tsvwg-l4s-arch/

[5] https://datatracker.ietf.org/doc/draft-ietf-tsvwg-nqb/

[6] *CableLabs DOCSIS 3.1 MAC and Upper Layer Protocols Interface Specification*, CM-SP-MULPIv3.1, I23, 2022

[7] https://datatracker.ietf.org/doc/pdf/draft-ietf-tsvwg-l4sops-03

[8] *A Latency Measurement System Using STAMP*, by Karthik Sundaresan, SCTE Expo 2021

[9] https://datatracker.ietf.org/meeting/114/materials/slides-114-tsvwg-update-on-l4s-work-in-ietf-114-hackathon-00

